



US009529773B2

(12) **United States Patent**
Hussain et al.

(10) **Patent No.:** **US 9,529,773 B2**
(45) **Date of Patent:** ***Dec. 27, 2016**

(54) **SYSTEMS AND METHODS FOR ENABLING ACCESS TO EXTENSIBLE REMOTE STORAGE OVER A NETWORK AS LOCAL STORAGE VIA A LOGICAL STORAGE CONTROLLER**

(71) Applicant: **CAVIUM, INC.**, San Jose, CA (US)

(72) Inventors: **Muhammad Raghieb Hussain**, Saratoga, CA (US); **Vishal Murgai**, Cupertino, CA (US); **Manojkumar Panicker**, Sunnyvale, CA (US); **Faisal Masood**, San Jose, CA (US); **Brian Folsom**, Northborough, MA (US); **Richard Eugene Kessler**, Northborough, MA (US)

(73) Assignee: **CAVIUM, INC.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **15/041,892**

(22) Filed: **Feb. 11, 2016**

(65) **Prior Publication Data**

US 2016/0162438 A1 Jun. 9, 2016

Related U.S. Application Data

(63) Continuation-in-part of application No. 14/300,552, filed on Jun. 10, 2014, now Pat. No. 9,294,567.
(Continued)

(51) **Int. Cl.**
G06F 3/06 (2006.01)
G06F 15/167 (2006.01)

(Continued)

(52) **U.S. Cl.**

CPC **G06F 15/167** (2013.01); **G06F 3/06** (2013.01); **G06F 3/0604** (2013.01); **G06F 3/065** (2013.01);

(Continued)

(58) **Field of Classification Search**

CPC G06F 3/0662; G06F 3/0665; G06F 3/067; G06F 3/0683; G06F 3/0685; G06F 3/0689; G06F 3/0626; G06F 3/0622; G06F 3/0617; G06F 3/0604; G06F 9/45558; G06F 9/45533; G06F 9/5072; G06F 13/1642; H04L 67/1097

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,329,318 A 7/1994 Keith
6,990,395 B2 1/2006 Ransom
(Continued)


Primary Examiner — Le H Luu

(74) *Attorney, Agent, or Firm* — Duane Morris LLP;
David T. Xue

(57) **ABSTRACT**


A new approach is proposed that contemplates systems and methods to support elastic (extensible/flexible) storage access in real time by mapping a plurality of remote storage devices that are accessible over a network fabric as logical namespace(s) via a logical storage controller using a multitude of access mechanisms and storage network protocols. The logical storage controller exports and presents the remote storage devices to one or more VMs running on a host of the logical storage controller as the logical namespace(s), wherein these remote storage devices appear virtually as one or more logical volumes of a collection of logical blocks in the logical namespace(s) to the VMs. As a result, each of the VMs running on the host can access these remote storage devices to perform read/write operations as if they were local storage devices via the logical namespace(s).


24 Claims, 6 Drawing Sheets



Related U.S. Application Data		2009/0037680 A1 *	2/2009	Colbert	G06F 3/0617
(60)	Provisional application No. 62/214,494, filed on Sep. 4, 2015.	2009/0307388 A1 *	12/2009	Tchapda	G06F 3/067 711/162 710/33
(51)	Int. Cl.	2010/0082922 A1	4/2010	George	
	H04L 12/26 (2006.01)	2010/0325371 A1	12/2010	Jagadish	
	H04L 29/08 (2006.01)	2012/0042033 A1	2/2012	Ayala, Jr.	
	G06F 9/455 (2006.01)	2012/0042034 A1	2/2012	Goggin	
	H04L 12/24 (2006.01)	2012/0110259 A1	5/2012	Mills	
(52)	U.S. Cl.	2012/0150933 A1	6/2012	Boersma	
	CPC	2013/0014103 A1	1/2013	Reuther	
	G06F 3/067 (2013.01); G06F 3/0611	2013/0042056 A1	2/2013	Shats	
	(2013.01); G06F 3/0662 (2013.01); G06F	2013/0097369 A1	4/2013	Talagala	
	3/0665 (2013.01); G06F 3/0683 (2013.01);	2013/0191590 A1	7/2013	Malwankar	
	G06F 3/0685 (2013.01); G06F 3/0689	2013/0198312 A1	8/2013	Tamir	
	(2013.01); G06F 9/45533 (2013.01); G06F	2013/0204849 A1 *	8/2013	Chacko	G06F 3/0604 707/692
	9/45558 (2013.01); H04L 41/0866 (2013.01);	2013/0318197 A1	11/2013	Plaisted	
	H04L 43/04 (2013.01); H04L 67/1097	2014/0007189 A1 *	1/2014	Huynh	G06F 3/0622 726/3
	(2013.01); G06F 2009/4557 (2013.01); G06F	2014/0059226 A1 *	2/2014	Messerli	G06F 9/5072 709/226
	2009/45583 (2013.01); G06F 2009/45587	2014/0089276 A1	3/2014	Satish	
	(2013.01)	2014/0173149 A1 *	6/2014	Walker	G06F 9/45533 710/263
(56)	References Cited	2014/0195634 A1	7/2014	Kishore	
	U.S. PATENT DOCUMENTS	2014/0281040 A1	9/2014	Liu	
	8,214,539 B1 7/2012 Kulanko	2014/0282521 A1 *	9/2014	Lango	G06F 9/45533 718/1
	8,239,655 B2 8/2012 Goggin	2014/0317206 A1 *	10/2014	Lomelino	H04L 67/1097 709/206
	8,291,135 B2 10/2012 Subramanian	2014/0331001 A1	11/2014	Liu	
	8,756,441 B1 6/2014 Mullins	2015/0120971 A1 *	4/2015	Bae	G06F 3/0626 710/106
	9,098,214 B1 8/2015 Vincent				
	2005/0060590 A1 3/2005 Bradley				
	2008/0235293 A1 9/2008 Levering				

* cited by examiner





PCIe/NVMe
Connection 211




FIG. 3

Namespace(s)/Logical volume(s)	Remote physical storage volumes
Ns_1	Vol_1
Ns_2	Vol_2
Ns_2	Vol_3
Ns_m	Vol_n

FIG. 4




FIG. 5




FIG. 6

1

SYSTEMS AND METHODS FOR ENABLING ACCESS TO EXTENSIBLE REMOTE STORAGE OVER A NETWORK AS LOCAL STORAGE VIA A LOGICAL STORAGE CONTROLLER

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part application of U.S. patent application Ser. No. 14/300,552, filed Jun. 10, 2014, which further claims the benefit of U.S. Provisional Patent Application No. 61/987,956, filed May 2, 2014 and is entitled "Systems and methods for enabling access to extensible storage devices over a network as local storage via NVMe controller," which is incorporated herein in its entirety by reference.

This application also claims the benefit of U.S. Provisional Patent Application No. 62/314,494, filed Sep. 4, 2015 and entitled "Enabling elastic storage over a network via a logical storage controller," which is incorporated herein in its entirety by reference.

BACKGROUND

Service providers have been increasingly providing their web services (e.g., web sites) at third party data centers in the cloud by running a plurality of virtual machines (VMs) on a host/server at the data center. Here, a VM is a software implementation of a physical machine (i.e. a computer) that executes programs to emulate an existing computing environment such as an operating system (OS). The VM runs on top of a hypervisor, which creates and runs one or more VMs on the host. The hypervisor presents each VM with a virtual operating platform and manages the execution of each VM on the host. By enabling multiple VMs having different operating systems to share the same host machine, the hypervisor leads to more efficient use of computing resources, both in terms of energy consumption and cost effectiveness, especially in a cloud computing environment.

Non-volatile memory express, also known as NVMe or NVM Express, is a specification that allows a solid-state drive (SSD) to make effective use of a high-speed Peripheral Component Interconnect Express (PCIe) bus attached to a computing device or host. Here the PCIe bus is a high-speed serial computer expansion bus designed to support hardware I/O virtualization and to enable maximum system bus throughput, low I/O pin count and small physical footprint for bus devices. NVMe typically operates on a non-volatile memory controller of the host, which manages the data stored on the non-volatile memory (e.g., SSD) and communicates with the host. Such an NVMe controller provides a command set and feature set for PCIe-based SSD access with the goals of increased and efficient performance and interoperability on a broad range of enterprise and client systems. The main benefits of using an NVMe controller to access PCIe-based SSDs are reduced latency, increased Input/Output (I/O) operations per second (IOPS) and lower power consumption, in comparison to Serial Attached SCSI (SAS)-based or Serial ATA (SATA)-based SSDs through the streamlining of the I/O stack.

Currently, a VM running on the host can access the PCIe-based SSDs via the physical NVMe controller attached to the host and the number of storage volumes the VM can access is constrained by the physical limitation on the maximum number of physical storage units/volumes that can be locally coupled to the physical NVMe controller.

2

Since the VMs running on the host at the data center may belong to different web service providers and each of the VMs may have its own storage needs that may change in real time during operation and are thus unknown to the host, it is impossible to predict and allocate a fixed amount of storage volumes ahead of time for all the VMs running on the host that will meet their storage needs. It is thus desirable to be able to provide storage volumes to the VMs that are extensible dynamically during real time operation via the NVMe controller.

The foregoing examples of the related art and limitations related therewith are intended to be illustrative and not exclusive. Other limitations of the related art will become apparent upon a reading of the specification and a study of the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Aspects of the present disclosure are best understood from the following detailed description when read with the accompanying figures. It is noted that, in accordance with the standard practice in the industry, various features are not drawn to scale. In fact, the dimensions of the various features may be arbitrarily increased or reduced for clarity of discussion.

FIG. 1 depicts an example of a diagram of a system to support virtualization of remote storage devices to be presented as local storage devices to VMs in accordance with some embodiments.

FIG. 2 depicts an example of a diagram of a system to support virtualization of remote storage devices to be presented as local storage devices to VMs via an NVMe controller in accordance with some embodiments.

FIG. 3 depicts an example of hardware implementation of the physical NVMe controller depicted in FIG. 2 in accordance with some embodiments.

FIG. 4 depicts a non-limiting example of a lookup table that maps between the NVMe namespaces of the logical volumes and the remote physical storage volumes in accordance with some embodiments.

FIG. 5 depicts a flowchart of an example of a process to support virtualization of remote storage devices to be presented as local storage devices to VMs in accordance with some embodiments.

FIG. 6 depicts a non-limiting example of a diagram of a system to support virtualization of a plurality of remote storage devices to be presented as local storage devices to VMs, wherein the physical NVMe controller further includes a plurality of virtual NVMe controllers in accordance with some embodiments.

DETAILED DESCRIPTION

The following disclosure provides many different embodiments, or examples, for implementing different features of the subject matter. Specific examples of components and arrangements are described below to simplify the present disclosure. These are, of course, merely examples and are not intended to be limiting. In addition, the present disclosure may repeat reference numerals and/or letters in the various examples. This repetition is for the purpose of simplicity and clarity and does not in itself dictate a relationship between the various embodiments and/or configurations discussed.

A new approach is proposed that contemplates systems and methods to support elastic (extensible/flexible) storage access in real time by mapping a plurality of remote storage

devices that are accessible over a network fabric as logical namespace(s) via a logical storage controller using a multitude of access mechanisms and storage network protocols. The logical storage controller exports and presents the remote storage devices to one or more VMs running on a host of the logical storage controller as the logical namespace(s), wherein these remote storage devices appear virtually as one or more logical volumes of a collection of logical blocks in the logical namespace(s) to the VMs. As a result, each of the VMs running on the host can access these remote storage devices to perform read/write operations as if they were local storage devices via the logical namespace(s).

By virtualizing and presenting the remote storage devices as if they were local disks to the VMs via the logical storage controller, the proposed elastic storage approach enables the VMs running on the server/host to access remote storage devices accessible over a network on demand in real time, removing any physical limitation on the number of storage volumes accessible by the VMs. Under such an approach, every storage device is presented to the VMs as local regardless of whether the storage device is locally attached or remotely accessible and the host of the VMs is enabled to dynamically allocate storage volumes to the VMs in real time based on their actual storage needs during operation instead of pre-allocating the storage volumes ahead of time, which may either be insufficient or result in unused storage space.

FIG. 1 depicts an example of a diagram of system 100 to support virtualization of remote storage devices to be presented as local storage devices to VMs. Although the diagrams depict components as functionally separate, such depiction is merely for illustrative purposes. It will be apparent that the components portrayed in this figure can be arbitrarily combined or divided into separate software, firmware and/or hardware components. Furthermore, it will also be apparent that such components, regardless of how they are combined or divided, can execute on the same host or multiple hosts, and wherein the multiple hosts can be connected by one or more networks.

In the example of FIG. 1, a computing unit/appliance/host 112 runs a plurality of VMs 110, each configured to provide a web-based service to clients over the Internet. Here, the host 112 can be a computing device, a communication device, a storage device, or any electronic device capable of running a software component. For non-limiting examples, a computing device can be, but is not limited to, a laptop PC, a desktop PC, a mobile device, or a server machine such as an x86/ARM server. A communication device can be, but is not limited to, a mobile phone.

In the example of FIG. 1, a logical or elastic storage controller 101 is configured to utilize a local network interface card/controller (NIC) 114 coupled to the host 112 to map one or more remote storage devices 122 accessible from a network fabric 132 via a network storage protocol and (optional) one or more locally coupled storage devices 120 accessible via a local bus to a plurality of logical volumes and present the logical volumes to the VMs 110 as if they were local logical storage devices. Here, the local NIC 114 is configured to include/implement a multitude of access mechanisms comprising but not limited to one or more of an NVMe controller 102, a Simplified Layer 2 Transport over Ethernet 103, a remote direct memory access (RDMA)/RDMA over Converged Ethernet (RoCE) controller 104, an internet Wide Area RDMA Protocol (iWARP) or Ceph controller 106, a SCSI or SATA controller 108.

In the example of FIG. 1, each of the locally coupled storage devices 120 and the remotely accessible storage devices 122 can be a non-volatile (non-transient) storage device, which can be but is not limited to, a solid-state drive (SSD), a Static random-access memory (SRAM), a magnetic hard disk drive, and a flash drive. The remote storage devices 122 can be part of one or more backend storage servers accessible by the NIC 114 over the network fabric 132.

In the example of FIG. 1, the network fabric 132 refers to switched fabric or switching fabric, which is a network topology in which network nodes interconnect via one or more network switches (e.g., crossbar switches). Because a switched fabric network spreads network traffic across multiple physical links, it yields higher total throughput than broadcast networks. The network fabric 132 constitutes not only the network (which can be but is not limited to internet, intranet, wide area network (WAN), local area network (LAN), wireless network, Bluetooth, WiFi, mobile communication network, or any other network type), but also the storage network protocols that establish the connectivity between one or more initiators (e.g., the logical storage controller 101/logical volumes) and targets (remote/backend storage devices) of the network.

FIG. 2 depicts an example of a diagram of system 200 to support virtualization of remote storage devices to be presented as local storage devices to VMs via an NVMe controller. Although the NVMe controller 102 is used as a non-limiting example to be utilized the logical storage controller to illustrate the proposed approach in the following discussions, a person ordinarily skilled in the art would understand that the same approach is also applicable to other types of logical storage controllers/technologies (e.g., Simplified Layer 2 Transport 103, RDMA/RoCE 104, iWARP 106, and SCSI/SATA 108) discussed above to achieve the same effects.

In the example of FIG. 2, the host 112 is coupled to the physical NVMe controller 102 via a PCIe/NVMe link/connection 211 and the VMs 110 running on the host 112 are configured to access the physical NVMe controller 102 via the PCIe/NVMe link/connection 211. For a non-limiting example, the PCIe/NVMe link/connection 211 is a PCIe Gen3 x8 bus.

In the example of FIG. 2, the physical NVMe controller 102 having at least an NVMe storage proxy engine 204, NVMe access engine 206 and a storage access engine 208 running on the NVMe controller 102. Here, the physical NVMe controller 102 is a hardware/firmware NVMe module having software, firmware, hardware, and/or other components that are used to effectuate a specific purpose. As discussed in details below, the physical NVMe controller 102 comprises one or more of a CPU or microprocessor, a storage unit or memory (also referred to as primary memory) such as RAM, with software instructions stored for practicing one or more processes. The physical NVMe controller 102 provides both Physical Functions (PFs) and Virtual Functions (VFs) to support the engines running on it, wherein the engines will typically include software instructions that are stored in the storage unit of the physical NVMe controller 102 for practicing one or more processes. As referred to herein, a PF function is a PCIe function used to configure and manage the single root I/O virtualization (SR-IOV) functionality of the controller such as enabling virtualization and exposing PCIe VFs, wherein a VF function is a lightweight PCIe function that supports SR-IOV and represents a virtualized instance of the controller 102. Each VF shares one or more physical resources on the physical

NVMe controller **102**, wherein such resources include but are not limited to on-controller memory **308**, hardware processor **306**, interface to storage devices **322**, and network driver **320** of the physical NVMe controller **102** as depicted in FIG. 3 and discussed in details below.

FIG. 3 depicts an example of hardware implementation **300** of the physical NVMe controller **102** depicted in FIG. 2. As shown in the example of FIG. 3, the hardware implementation **300** includes at least an NVMe processing engine **302**, and an NVMe Queue Manager (NQM) **304** implemented to support the NVMe processing engine **302**. Here, the NVMe processing engine **302** includes one or more CPUs/processors **306** (e.g., a multi-core/multi-threaded ARM/MIPS processor), and a primary memory **308** such as DRAM. The NVMe processing engine **302** is configured to execute all NVMe instructions/commands and to provide results upon completion of the instructions. The hardware-implemented NQM **304** provides a front-end interface to the engines that execute on the NVMe processing engine **302**. In some embodiments, the NQM **304** manages at least a submission queue **312** that includes a plurality of administration and control instructions to be processed by the NVMe processing engine **302** and a completion queue **314** that includes status of the plurality of administration and control instructions that have been processed by the NVMe processing engine **302**. In some embodiments, the NQM **304** further manages one or more data buffers **316** that include data read from or to be written to a storage device via the NVMe controllers **102**. In some embodiments, one or more of the submission queue **312**, completion queue **314**, and data buffers **316** are maintained within memory **310** of the host **112**. In some embodiments, the hardware implementation **300** of the physical NVMe controller **102** further includes an interface to storage devices **222**, which enables the plurality of storage devices **120** to be coupled to and accessed by the physical NVMe controller **102** locally via a local bus, and a network driver **230**, which enables a plurality of storage devices **122** to be connected to the NVMe controller **102** remotely over the network fabric **132** via the multitude of access mechanisms discussed above.

In the example of FIG. 2, the NVMe access engine **206** of the NVMe controller **102** is configured to receive and manage instructions and data for read/write operations from the VMs **110** running on the host **112**. When one of the VMs **110** running on the host **112** performs a read or write operation, it places a corresponding instruction in a submission queue **312**, wherein the instruction is in NVMe format. During its operation, the NVMe access engine **206** utilizes the NQM **304** to fetch the administration and/or control commands from the submission queue **312** on the host **112** based on a "doorbell" of read or write operation, wherein the doorbell is generated by the VM **110** and received from the host **112**. The NVMe access engine **206** also utilizes the NQM **304** to fetch the data to be written by the write operation from one of the data buffers **316** on the host **112**. The NVMe access engine **206** then places the fetched commands in a waiting buffer **318** in the memory **308** of the NVMe processing engine **302** waiting for the NVMe Storage Proxy Engine **204** to process. Once the instructions are processed, The NVMe access engine **206** puts the status of the instructions back in the completion queue **314** and notifies the corresponding VM **110** accordingly. The NVMe access engine **206** also puts the data read by the read operation to the data buffer **316** and makes it available to the VM **110**.

In some embodiments, each of the VMs **110** running on the host **112** has an NVMe driver **214** configured to interact with the NVMe access engine **206** of the NVMe controller **102** via the PCIe/NVMe link/connection **211**. In some embodiments, each of the NVMe driver **214** is a virtual function (VF) driver configured to interact with the PCIe/NVMe link/connection **211** of the host **112** and to set up a communication path between its corresponding VM **110** and the NVMe access engine **206** and to receive and transmit data associated with the corresponding VM **110**. In some embodiments, the VF NVMe driver **214** of the VM **110** and the NVMe access engine **206** communicate with each other through a SR-IOV PCIe connection **211** as discussed above.

In some embodiments, the VMs **110** run independently on the host **112** and are isolated from each other so that one VM **110** cannot access the data and/or communication of any other VMs **110** running on the same host. When transmitting commands and/or data to and/or from a VM **110**, the corresponding VF NVMe driver **214** directly puts and/or retrieves the commands and/or data from its queues and/or the data buffer, which is sent out or received from the NVMe access engine **206** without the data being accessed by the host **112** or any other VMs **110** running on the same host **112**.

In the example of FIG. 2, the storage access engine **208** of the NVMe controller **102** is configured to access and communicate with a plurality of non-volatile disk storage devices/units, wherein each of the storage units is either (optionally) locally coupled to the NVMe controller **102** via the interface to storage devices **222** (e.g., local storage devices **120**), or remotely accessible by the physical NVMe controller **102** over the network fabric **132** (e.g., remote storage devices **122**) via the network communication interface/driver **230** following certain communication protocols such as TCP/IP protocol.

In the example of FIG. 2, the NVMe storage proxy engine **204** of the NVMe controller **102** is configured to collect volumes of the remote storage devices accessible via the storage access engine **208** over the network under the storage network protocol and convert the storage volumes of the remote storage devices to one or more NVMe namespaces each including a plurality of logical volumes (a collection of logical blocks) to be accessed by VMs **110** running on the host **112**. As such, the NVMe namespaces may cover both the storage devices locally attached to the NVMe controller **102** and those remotely accessible by the storage access engine **208** under the storage network protocol. The storage network protocol is used to access a remote storage device accessible over the network fabric **132**, wherein such storage network protocol can be but is not limited to Internet Small Computer System Interface (iSCSI). iSCSI is an Internet Protocol (IP)-based storage networking standard for linking data storage devices by carrying SCSI commands over the networks. By enabling access to remote storage devices over the network, iSCSI increases the capabilities and performance of storage data transmission over local area networks (LANs), wide area networks (WANs), and the Internet.

In some embodiments, the NVMe storage proxy engine **204** organizes and maps the remote storage devices to one or more logical or virtual volumes/blocks in the NVMe namespaces, to which the VMs **110** can access and perform I/O operations as if they were local storage volumes. Here, each volume is classified as logical or virtual since it maps to one or more physical storage devices either locally attached to or remotely accessible by the NVMe controller **102** via the storage access engine **208**. In some embodi-

ments, multiple VMs **110** running on the host **112** are enabled to access the same logical volume or virtual volume and each logical/virtual volume can be shared among multiple VMs. In some embodiments, the virtual volume includes a meta-data mapping table between the logical volume and the remote storage devices **122**, wherein the mapping table translates an incoming (virtual) volume identifier and a logical block addressing (LBA) on the virtual volume to one or more corresponding physical disk identifiers and LBAs on the storage devices. In some embodiments, the logical volume may include logical blocks across multiple physical disks in the storage devices. For non-limiting examples, two possible scenarios are:

A single logical volume could be mapped to and served by a single remote storage device/backend server using a given underlying technology, e.g., NVMe, Simplified Transport, iWARP, RoCE, etc.

A single logical volume could be mapped to and served by a plurality of remote storage devices under one or more different controllers/technologies via the NIC **114**. For a non-limiting example, the logical storage controller **101** may utilize the NVMe controller **102** to map from one remote server, and use RDMA/RoCE controller **104** to map from the other remote server, wherein the logical storage controller **101** is configured to handle the logic and complexity among the different underlying controllers and technologies.

In some embodiments, the NVMe storage proxy engine **204** establishes a lookup table that maps between the NVMe namespaces of the logical volumes, $Ns_1, \dots, Ns_m$, and the remote physical storage devices/volumes, $Vol_1, \dots, Vol_n$, accessible over the network as shown by the non-limiting example depicted in FIG. 4. Here, there is a multiple-to-multiple correspondence between the NVMe namespaces and the physical storage volumes, meaning that one namespace (e.g., Ns_2) may correspond to a logical volume that maps to a plurality of remote physical storage volumes (e.g., Vol_2 and Vol_3), and a single remote physical storage volume may also be included in a plurality of logical volumes and accessible by the VMs **110** via their corresponding NVMe namespaces. In some embodiments, the NVMe storage proxy engine **204** is configured to expand the mappings between the NVMe namespaces of the logical volumes and the remote physical storage devices/volumes to add additional storage volumes on demand. Specifically, when at least one of the VMs **110** running on the host **112** requests for more storage volumes, the NVMe storage proxy engine **204** may expand the namespace/logical volume accessed by the VM to include additional remote physical storage devices.

For a non-limiting example, a VM **110** running on the host **112** may request the logical storage controller **101** in FIG. 1 to allocate X TB amount of disk space at runtime. The logical storage controller **101** is configured to utilize the NIC **114** (e.g., NVMe controller **102** depicted in FIG. 2) to allocate the required disk space from the remote storage devices **122** on the backend storage servers over the network fabric **132** as logical volumes to the VM **110**. Here, the logical storage controller **101** is configured to provision the requested storage space on the locally attached storage device **120** and the remote storage devices **122** via one or more provisioning mechanisms, which include but are not limited to, the NVMe controller **102**, one of RDMA/RoCE/iWARP controllers, or a Simplified Layer 2 Transport over the plain Ethernet fabric. It is up to the logical storage controller **101** to determine which controllers and/or technologies of the NIC **114** to utilize to accomplish the storage

request (and any addition request) by the VM **110**. If the VM **110** needs more storage at runtime, it may request for, for a non-limiting example, an additional storage capacity of Y TB, and the logical storage controller **101** is configured to dynamically provision that increased storage capacity in real time via the NIC **114** to make X+Y TB amount of storage available to the VM **110**. The logical storage controller **101** is configured to accomplish this by either extending the storage of the already allocated logical volume of X TB by Y TB or by allocating another logical volume of X+Y TB of storage space. The logical storage controller **101** is further configured to expand mappings between the logical volumes and the remote physical storage devices to include the additional storage space allocated on the remote physical storage devices.

In some embodiments, the NVMe storage proxy engine **204** further includes an adaptation layer/shim **216**, which is a software component configured to manage message flows between the NVMe namespaces and the remote physical storage volumes. Specifically, when instructions for storage operations (e.g., read/write operations) on one or more logical volumes/namespaces are received from the VMs **110** via the NVMe access engine **206**, the adaptation layer/shim **216** converts the instructions under NVMe specification to one or more corresponding instructions on the remote physical storage volumes under the storage network protocol such as iSCSI according to the lookup table. Conversely, when results and/or feedbacks on the storage operations performed on the remote physical storage volumes are received via the storage access engine **208**, the adaptation layer/shim **216** also converts the results to feedbacks about the operations on the one or more logical volumes/namespaces and provides such converted results to the VMs **110**.

In the example of FIG. 2, the NVMe access engine **206** of the NVMe controller **102** is configured to export and present the NVMe namespaces and logical volumes of the remote physical storage devices **122** to the VMs **110** running on the host **112** as accessible storage devices that are no different from those locally connected storage devices **120**. The actual mapping, expansion, and operations on the remote storage devices **122** over the network using iSCSI-like storage network protocol performed by the NVMe controller **102** are transparent to the VMs **110**, which provides the instructions on the logical volumes that map to the remote storage volumes.

FIG. 5 depicts a flowchart of an example of a process to support virtualization of remote storage devices to be presented as local storage devices to VMs. Although this figure depicts functional steps in a particular order for purposes of illustration, the process is not limited to any particular order or arrangement of steps. One skilled in the relevant art will appreciate that the various steps portrayed in this figure could be omitted, rearranged, combined and/or adapted in various ways.

In the example of FIG. 5, the flowchart **500** starts at block **502**, where one or more logical volumes in one or more NVMe namespaces are created and mapped to a plurality of remote storage devices accessible over a network. The flowchart **500** continues to block **504**, where the NVMe namespaces of the logical volumes are presented to one or more virtual machines (VMs) running on a host as if they were local storage volumes. The flowchart **500** continues to block **506**, where a first instruction to perform a read/write operation on one of the logical volumes mapped to the remote storage devices is received from one of the VMs. The flowchart **500** continues to block **508**, where the namespaces of the logical volumes in the first instruction is converted to

storage volumes of the remote storage devices in a second instruction according to a storage network protocol. The flowchart 500 ends at block 410, where result and/or data of the read/write operation performed on the remote storage devices is received, processed and presented to the VM after the read/write operation is performed on the storage volumes of the remote storage devices over the network using the second instruction.

FIG. 6 depicts a non-limiting example of a diagram of system 600 to support virtualization of remote storage devices as local storage devices for VMs, wherein the physical NVMe controller 102 further includes a plurality of virtual NVMe controllers 602. In the example of FIG. 6, the plurality of virtual NVMe controllers 602 run on the single physical NVMe controller 102 where each of the virtual NVMe controllers 602 is a hardware accelerated software engine emulating the functionalities of an NVMe controller to be accessed by one of the VMs 110 running on the host 112. In some embodiments, the virtual NVMe controllers 602 have a one-to-one correspondence with the VMs 110, wherein each virtual NVMe controller 602 interacts with and allows access from only one of the VMs 110. Each virtual NVMe controller 602 is assigned to and dedicated to support one and only one of the VMs 110 to access its storage devices, wherein any single virtual NVMe controller 602 is not shared across multiple VMs 110.

In some embodiments, each virtual NVMe controller 602 is configured to support identity-based authentication and access from its corresponding VM 110 for its operations, wherein each identity permits a different set of API calls for different types of commands/instructions used to create, initialize and manage the virtual NVMe controller 602, and/or provide access to the logic volume for the VM 110. In some embodiments, the types of commands made available by the virtual NVMe controller 602 vary based on the type of user requesting access through the VM 110 and some API calls do not require any user login. For a non-limiting example, different types of commands can be utilized to initialize and manage virtual NVMe controller 602 running on the physical NVMe controller 102.

As shown in the example of FIG. 6, each virtual NVMe controller 602 may further include a virtual NVMe storage proxy engine 604 and a virtual NVMe access engine 606, which function in a similar fashion as the respective NVMe storage proxy engine 204 and a NVMe access engine 206 discussed above. In some embodiments, the virtual NVMe storage proxy engine 604 in each virtual NVMe controller 602 is configured to access both the locally attached storage devices 120 and remotely accessible storage devices 122 via the storage access engine 208, which can be shared by all the virtual NVMe controllers 602 running on the physical NVMe controller 102.

During operation, each virtual NVMe controller 602 creates and maps one or more logical volumes in one or more NVMe namespaces mapped to a plurality of remote storage devices accessible over a network. Each virtual NVMe controller 602 then presents the NVMe namespaces of the logical volumes to its corresponding VM 110 as if they were local storage volumes. When a first instruction to perform a read/write operation on the logical volumes is received from the VM 110, the virtual NVMe controller 602 converts the NVMe namespaces of the logical volumes in the first instruction to storage volumes of the remote storage devices in a second instruction according to a storage network protocol. The virtual NVMe controller 602 presents result and/or data of the read/write operation to the VM 110 after the read/write operation has been performed on the

storage volumes of the remote storage devices over the network using the second instruction.

In some embodiments, each virtual NVMe controller 602 depicted in FIG. 6 has one or more pairs of submission queue 312 and completion queue 314 associated with it, wherein each queue can accommodate a plurality of entries of instructions from one of the VMs 110. As discussed above, the instructions in the submission queue 312 are first fetched by the NQM 304 from the memory 310 of the host 112 to the waiting buffer 318 of the NVMe processing engine 302 as discussed above. During its operation, each virtual NVMe controller 602 retrieves the instructions from its corresponding VM 110 from the waiting buffer 318 and converts the instructions according to the storage network protocol in order to perform a read/write operation on the data stored on the local storage devices 120 and/or remote storage devices 122 over the network by invoking VF functions provided by the physical NVMe controller 102. During the operation, data is transmitted to or received from the local/remote storage devices in the logical volume of the VM 110 via the interface to storage access engine 208. Once the operation has been processed, the virtual NVMe controller 602 saves the status of the executed instructions in the waiting buffer 318 of the processing engine 302, which are then placed into the completion queue 314 by the NQM 304. The data being processed by the instructions of the VMs 110 is also transferred between the data buffer 316 of the memory 310 of the host 112 and the memory 308 of the NVMe processing engine 302.

The methods and system described herein may be at least partially embodied in the form of computer-implemented processes and apparatus for practicing those processes. The disclosed methods may also be at least partially embodied in the form of tangible, non-transitory machine readable storage media encoded with computer program code. The media may include, for example, RAMs, ROMs, CD-ROMs, DVD-ROMs, BD-ROMs, hard disk drives, flash memories, or any other non-transitory machine-readable storage medium, wherein, when the computer program code is loaded into and executed by a computer, the computer becomes an apparatus for practicing the method. The methods may also be at least partially embodied in the form of a computer into which computer program code is loaded and/or executed, such that, the computer becomes a special purpose computer for practicing the methods. When implemented on a general-purpose processor, the computer program code segments configure the processor to create specific logic circuits. The methods may alternatively be at least partially embodied in a digital signal processor formed of application specific integrated circuits for performing the methods.

The foregoing description of various embodiments of the claimed subject matter has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the claimed subject matter to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art. Embodiments were chosen and described in order to best describe the principles of the invention and its practical application, thereby enabling others skilled in the relevant art to understand the claimed subject matter, the various embodiments and with various modifications that are suited to the particular use contemplated.

11

What is claimed is:

1. A system to support elastic network storage, comprising:

a logical storage controller of a local network interface card/controller (NIC), configured to:

- accept a request for storage space from one of a plurality of virtual machines (VMs) running on a host;
- allocate storage volumes on one or more remote storage devices accessible over a network fabric in accordance with the request for the storage space;
- create and map one or more logical volumes in one or more logical namespaces to the storage volumes on the remote storage devices;
- present the logical volumes mapped to the storage volumes on the remote storage devices as local storage volumes to the VM requesting the storage space;
- enable the VM to perform a read/write operation on the logical volumes via a first instruction.

2. The system of claim 1, wherein:

the host of the VM is an x86 or ARM server.

3. The system of claim 1, wherein:

the NIC includes or implements one or more of an NVMe controller, a remote direct memory access (RDMA)/RDMA over Converged Ethernet (RoCE) controller, a Simplified Transport over Ethernet, an internet Wide Area RDMA Protocol (iWARP) controller, and a local SCSI or SATA controller.

4. The system of claim 1, wherein:

the remote storage devices are part of one or more backend storage servers accessible by the NIC over the network fabric.

5. The system of claim 1, wherein:

the network fabric is a network topology in which nodes in a network interconnect via one or more network switches.

6. The system of claim 1, wherein:

the logical volumes are further mapped to a storage device attached to the NIC locally.

7. The system of claim 1, wherein:

the VMs run independently on the host and are isolated from each other so that one VM cannot access the data and/or communication of any other VMs running on the same host.

8. The system of claim 1, wherein:

the logical storage controller is configured to

- convert the logical volumes in the first instruction from the VM to the storage volumes of the remote storage devices in a second instruction according to a storage network protocol;
- present result and/or data of the read/write operation to the VM after the read/write operation has been performed on the storage volumes of the remote storage devices over the network fabric using the second instruction.

9. The system of claim 8, wherein:

the storage network protocol is Internet Small Computer System Interface (iSCSI).

10. The system of claim 1, wherein:

the logical storage controller is configured to enable multiple of the plurality of VMs to access the same logical volume and each logical volume is enabled to be shared among the multiple VMs.

12

11. The system of claim 1, wherein:

the logical storage controller is configured to map a single logical volume to a single remote storage device via a controller of the NIC.

12. The system of claim 1, wherein:

the logical storage controller is configured to map a single logical volume to a plurality of remote storage devices via different controllers of the NIC.

13. The system of claim 1, wherein:

the logical storage controller is configured to perform the mapping and the operations on the remote storage devices over the network transparent to the VM.

14. The system of claim 1, wherein:

the logical storage controller is configured to

- dynamically allocate additional storage space on the remote physical storage devices to the VM per additional storage request by the VM in real time;
- expand mappings between the logical volumes and the remote physical storage devices to include the additional storage space allocated on the remote physical storage devices.

15. A method to support elastic network storage via a local network interface card/controller (NIC), comprising:

- accepting a request for storage space from one of a plurality of virtual machines (VMs) running on a host;
- allocating storage volumes on one or more remote storage devices accessible over a network fabric in accordance with the request for the storage space;

- creating and mapping one or more logical volumes in one or more logical namespaces to the storage volumes on the remote storage devices;

- presenting the logical volumes mapped to the storage volumes on the remote storage devices as local storage volumes to the VM requesting the storage space;

- enabling the VM to perform a read/write operation on the logical volumes via a first instruction.

16. The method of claim 15, further comprising:

- including or implementing one or more of an NVMe controller, a remote direct memory access (RDMA)/RDMA over Converged Ethernet (RoCE) controller, a Simplified Transport over Ethernet, an internet Wide Area RDMA Protocol (iWARP) controller, and a local SCSI or SATA controller via the NIC.

17. The method of claim 15, further comprising:

- mapping the logical volumes to a storage device attached to the NIC locally.

18. The method of claim 15, further comprising:

- enabling the VMs to run independently on the host and isolating the VMs from each other so that one VM cannot access the data and/or communication of any other VMs running on the same host.

19. The method of claim 15, further comprising:

- converting the logical volumes in the first instruction from the VM to the storage volumes of the remote storage devices in a second instruction according to a storage network protocol;

- presenting result and/or data of the read/write operation to the VM after the read/write operation has been performed on the storage volumes of the remote storage devices over the network fabric using the second instruction.

20. The method of claim 15, further comprising:

- enabling multiple of the plurality of VMs to access the same logical volume and each logical volume to be shared among the multiple VMs.

21. The method of claim 15, further comprising:

- mapping a single logical volume to a single remote storage device via a controller of the NIC.

13

14

22. The method of claim 15, further comprising:
mapping a single logical volume to a plurality of remote
storage devices via different controllers of the NIC.

23. The method of claim 15, further comprising:
performing the mapping and the operations on the remote 5
storage devices over the network transparent to the
VM.

24. The method of claim 15, further comprising:
dynamically allocating additional storage space on the
remote physical storage devices to the VM per addi- 10
tional storage request by the VM in real time;
expanding mappings between the logical volumes and the
remote physical storage devices to include the addi-
tional storage space allocated on the remote physical
storage devices. 15

* * * * *